

CLASSICAL AND QUANTUM REALTIME ALTERNATING AUTOMATA

H. Gökalp Demirci^(C) Mika Hirvensalo^(A)
 Klaus Reinhardt^(D) A. C. Cem Say^(E)
 Abuzer Yakaryılmaz^(B)

^(C)University of Chicago, Department of Computer Science, Chicago, IL 60637 USA
 demirci@cs.uchicago.edu

^(A)Department of Mathematics and Statistics, University of Turku, FI-20014 Turku Finland
 mikhirve@utu.fi

^(D)Wilhelm-Schickard-Institut für Informatik, University of Tübingen, Germany
 currently Humboldt University of Berlin, Germany
 klaus.reinhardt@uni-tuebingen.de

^(E)Boğaziçi University, Department of Computer Engineering, Bebek 34342 İstanbul, Turkey
 say@boun.edu.tr

^(B)National Laboratory for Scientific Computing, Petrópolis, RJ, 25651-075, Brazil
 abuzer@lncc.br

Abstract

We present some new results on realtime classical and quantum alternating models. Firstly, we show that the emptiness problem for alternating one-counter automata on unary alphabets is undecidable. Then, we define realtime private alternating finite automata (PAFAs) and show that they can recognize some non-regular unary languages, and the emptiness problem is undecidable for them. Moreover, PAFAs augmented with a counter can recognize the unary squares language, which seems to be difficult even for some classical counter automata with two-way input. For quantum finite automata (QFAs), we show that the emptiness problem for universal QFAs on general alphabets and alternating QFAs with two alternations on unary alphabets are undecidable. On the other hand, the same problem is decidable for nondeterministic QFAs on general alphabets. We also show that the unary squares language is recognized by alternating QFAs with two alternations.

Keywords: *alternating finite automata, private and quantum alternation, realtime computation, counter automata, emptiness problem, unary languages*

^(A)Hirvensalo was partially supported by Väisälä Foundation grant 2011-2013.

^(B)Yakaryılmaz was partially supported by CAPES, ERC Advanced Grant MQC, and FP7 FET project QALGO.

1. Introduction

Alternation is a generalization of nondeterminism [4]. Although alternating finite automata (AFAs) can still recognize only the regular languages, even with a two-way input head, they can be more powerful when given more resources. For example, AFAs augmented with a counter (A1CAs) can recognize some unary non-regular languages [19], whereas their nondeterministic counterparts cannot recognize any unary non-regular language, even when allowed to pause the input head indefinitely on a symbol, and given a stack instead of a counter, thereby upgraded to a one-way nondeterministic pushdown automaton (PDA) [8].

It is a well-known fact that the emptiness problem for one-way PDAs is decidable, but the universality problem for them is undecidable [21]. Therefore, the emptiness problem is also undecidable for universal pushdown automata, and the same result can be shown for the real-time version by replacing the stack with a counter, obtaining a A1CA with a single alternation having only universal states. But, the answer is not trivial for A1CAs with unary input alphabets. In this paper, we prove that the emptiness problem for A1CAs on unary alphabets is undecidable.

Private alternation is a generalization of alternation [17, 16, 18], which is usually modelled as a game between the existential and universal players (see also [5]), where some computational resources (tape heads, working memories, etc) are private to the universal player. Such privacy can increase the computational power of models. For example, if the input head is private, then one-way private AFAs can recognize any alternating linear-space language ($\text{ASPACE}(n) = \text{DTIME}(2^n)$) by simulating linear-space alternating Turing machines (TMs) on the given inputs. Note that the automaton runs forever in some useless computational paths during this simulation. In this paper, we define realtime versions of private AFAs, and show that (i) they can recognize some nonregular unary languages, (ii) the emptiness problem is undecidable for them, and (iii) their one-counter counterparts can recognize the unary squares language, which seems to be a hard language even for some classical two-way counter finite automata [10, 15, 2].

The concept of quantum alternation was introduced [24] and shown to be very powerful recently: One-way alternating quantum finite automata (QFAs) can simulate the computation of any TM on a given input, and so they recognize any recursively enumerable language. Similar to private AFAs, the automaton runs forever in some useless computational paths during this simulation. In this paper, we focus on their realtime counterparts and show that the emptiness problem is undecidable for both universal QFAs on general alphabets and alternating QFAs with two alternations on unary alphabet, but decidable for nondeterministic QFAs. Moreover, we show that the unary squares language is recognized by alternating QFAs with two alternations.

Throughout the paper, Σ denotes the input alphabet not containing the end-marker ($\$$) and $\tilde{\Sigma} = \Sigma \cup \{\$$. Any given input $w \in \Sigma$ is read as $\$w\$$. All realtime models can spent at most two steps on each input symbol so that they can make both existential and universal choices on the same symbol.

We classify the results with respect to models; alternation, private alternation, and quantum

alternation. We devote a section to each of them, and provide the required background within the sections.

2. Alternation

A realtime alternating finite automaton is (AFA) a 5-tuple

$$A = (S, \Sigma, \delta, s_1, s_a),$$

where

S is the finite set of states composed by existential and universal states,

s_1 is the initial state and s_a is the accepting state, and,

$\delta : S \times \tilde{\Sigma} \rightarrow \mathcal{P}(S)$ is the transition function,

where \mathcal{P} is the powerset of a given set. The automaton spends two steps on each symbol. When A is in state $s \in S$ and reads $\sigma \in \tilde{\Sigma}$, it switches to a set of states $S' \subseteq S$, where $\delta(s, \sigma) = S'$. The automaton gives the decision of acceptance if it ends the computation in state s_a (after reading ϕ). The computation of an AFA can be shown as a tree such that each node represents a state, and the edges represent the transitions. We associate each inner node with either a “ \vee ” or a “ \wedge ”, depending on whether its corresponding state is existential or universal, respectively. When evaluating the tree, each leaf in which the input is accepted (rejected) takes the value of “true” (“false”). The value of root can be evaluated from bottom to top: Any existential inner node (assigned a “ \vee ”) takes the value of “true” if at least one of its children has taken the value of “true” and any universal inner node (assigned a “ \wedge ”) takes the value of “true” if all of its children have taken the value of “true”. The input is accepted if and only if the root takes the value of “true”.

A realtime alternating one-counter automaton (A1CA) is a (realtime) AFA augmented with an integer counter. In each step, the automaton additionally can test the counter whether it is zero or not, and then it updates the counter by a value from $\{-1, 0, 1\}$ in each branch. That is, the transition function is formally defined as $\delta : S \times \tilde{\Sigma} \times \{0, \pm\} \rightarrow \mathcal{P}(S \times \{-1, 0, 1\})$, where 0 and \pm represent the value of the counter is zero or nonzero, respectively.

Theorem 2.1 *The emptiness problem for unary A1CAs is undecidable.*

Proof. For a deterministic Turing machine M , we construct a realtime alternating counter automaton A which accepts the unary input u^{2n} if and only if M starting on an empty tape halts in exactly n steps. To solve the emptiness problem for A is then to solve whether M halts, a known undecidable problem. The proof idea is similar to Theorem 3.4. in [4] showing how an alternating Turing machine can mimic deterministic $T(n)$ time computation in $O(\log T(n))$ space by utilizing backwards simulation.

Without loss of generality, we can assume that the tape of M is semi-infinite, cells numbered with nonnegative integers, that there is a special starting symbol \triangleright always in cell number 0,

and the machine never attempts to overwrite \triangleright , or move the head to the left of it. Moreover, by introducing another left endmarker, we may assume that the final step (if M halts) is the first time M returns to head position $C = 0$, and halts in a unique halting state q_f .

The simulation uses the input word u^{2n} as a clock measuring the steps of computation, and the counter to store the read-write head location of the simulated machine. By the aforementioned assumptions, the counter value will start at $C = 0$, remain nonnegative, and return to value $C = 0$ only if the computation halts. The starting configuration of M is numbered as 0, final as n , and the backwards simulation means that the configuration number $n - i$ is simulated when reading symbols $2i$ and $2i + 1$ of the input word u^{2n} .

We demonstrate that the simulation can be done with an alternating machine A having a finite memory and one integer-valued counter, and leave it to the reader to verify that A can be actually formalized as an alternating automaton with one counter. Since M is deterministic, the contents in position C of the t th configuration depends only on the contents in positions $C - 1$, C , and $C + 1$ of the $(t - 1)$ th configuration. Contents may be just a symbol of the alphabet of M or a state-symbol pair from M for the case that M is in the state and the input head is on this position (the later can only be the case for one of the three positions). Hence there is a finite partial function $s_C^{(t+1)} = \text{Next}_M(s_{C-1}^{(t)}, s_C^{(t)}, s_{C+1}^{(t)})$, where the superscript refers to configuration number, and the subscript to the position of the contents.

The machine A begins in a state representing the final Turing machine configuration: State q_f paired with the scanned symbol \triangleright in memory, and counter 0 representing this final head position. The task of machine A is to check whether such a configuration can be obtained by a computation of Turing machine M in n steps.

In each stage of the checking, machine A keeps the configuration contents s in its memory, and its target is to check whether $s = s_C^{(t)}$. It reads one input symbol and (existentially) guesses a combination (s_{-1}, s_0, s_1) with $s = \text{Next}_M(q, s_{-1}, s_0, s_1)$. Then A reads another input symbol and universally branches into configurations where C is replaced with $C - 1$, C , and $C + 1$, respectively, and enters again into the checking stage to verify whether $s_{C+d}^{(t-1)} = s_d$ for $d \in \{-1, 0, 1\}$.

It should be noted that the superindex t is not stored in the memory of the machine A , but it is not necessary: Machine A just stops checking when the input word u^{2n} is completely read, and accepts only if it is able to reach the initial configuration of M . \square

3. Private alternation

Private alternation [17] was introduced to model games where the universal player is able to hide some information from the existential player. In the original definition of private alternation [17, 16], the input head is considered to be a public resource, visible to both existential and universal states. Our definition of private alternation refers to the model in [6], where the input head is a resource private to the universal states, (i.e. its position cannot be seen by the

existential states).

A *realtime private alternating finite automaton* (PAFA) is a 12-tuple

$$M = (S_c, S_p, \Sigma, U, \delta_E, \delta_U, s_{c0}, s_{p0}, q_a, q_r, \Gamma, \Delta),$$

where

S_c is the finite set of common state components,

S_p is the finite set of private state components,

Σ is the input tape alphabet, (with $\epsilon \notin \Sigma$ as the input end-marker),

$U \subseteq (S_c \times S_p)$ is the set of universal states,

$\delta_E : S_c \rightarrow \mathcal{P}(S_c)$ is the existential transition function,

$\delta_U : S_c \times S_p \times (\Sigma \cup \{\epsilon\}) \rightarrow \mathcal{P}(S_c \times S_p)$ is the universal transition function,

$s_{c0} \in S_c$, and $s_{p0} \in S_p$ are the common and private components of the initial state, respectively,

$q_a, q_r \in S_c$ are the accepting and rejecting common state components,

and Γ and Δ , the public and private game alphabets, respectively, each containing at least two symbols, and satisfying $\Gamma \cap \Delta = \emptyset$, will be explained shortly.

Every state of such a PAFA is a pair from $S_c \times S_p$. $E = S_c \times S_p - U$ is the set of existential states. For transitions from U (resp., E), δ_U (resp., δ_E) is applied. The PAFA makes two transitions on each input symbol (or end-marker). In other words, the input head moves one step to the right after every two transitions. Existential moves (transitions from existential states) do not see and cannot change the private component of the state. They also do not see the current input symbol. A universal move (transitions from universal state), according to δ_U , depends on both the common and private state components, and can change both components.

For an input string w of length n , the computation of a PAFA on the input tape containing $\epsilon w \epsilon$ can be visualized as a *computation tree* of depth $2n + 4$, whose root (at level 0) is labeled with the initial state (s_{c0}, s_{p0}) , the children of a node $N = (s_c, s_p)$ at some level i correspond to the members of $\{(s'_c, s_p) \mid s'_c \in \delta_E(s_c)\}$ if N is existential, and to the members of $\{(s'_c, s'_p) \mid (s'_c, s'_p) \in \delta_U(s_c, s_p, \sigma)\}$, where σ is the $(\lceil \frac{i}{2} \rceil)$ th tape symbol, if N is universal. In addition to the usual AFA acceptance criterion that this tree should have a subtree containing the root, exactly one child of every existential node in all levels, and all children of every universal node, with all leaf nodes corresponding to states with accepting common state components, a PAFA imposes an additional condition on accepting computation subtrees, as we describe below.

The transition functions δ_E and δ_U were defined to return sets of “next” states. We impose additional conditions on the cardinalities of those sets, namely, δ_E either returns a singleton, or a set of size $|\Gamma|$, whereas δ_U either returns a singleton, or a set of size $|\Gamma \cup \Delta|$. We further require that each outgoing transition from a state S with such multiple successors is associated with a distinct member of Γ , if S is existential, and of $\Gamma \cup \Delta$, if S is universal.

The rationale for this addition to the model becomes evident when one views the computation of a PAFA as a game between two players, with the automaton itself playing the role of a game

board. As is the case with the analogous interpretation for AFAs, the existential player wants to lead the machine to accept the input string w , whereas the universal one aims rejection. Unlike the AFA scenario, however, the existential player can not see all moves of the universal player, and therefore has only partial knowledge of the machine state. Note that both players are assumed to know the string w in its entirety, and the fact that the existential player cannot see the currently scanned input symbol is a result of the input head being a private resource. States with a single outgoing transition do not require a choice of move by the players, and do not count as moves in this game. When a player does have a choice of move, each of its alternative moves has a name indicated by the corresponding symbol from the game alphabet. If the universal player wishes to do so, it may select a public move, i.e. one corresponding to a member of the public game alphabet Γ . Such moves must not make any changes in the current private state component. In this case, the existential player can see what its opponent has done. The power of private alternation ensues from the ability of the universal player to make private moves, corresponding to members of Δ , and change the private state component accordingly. The existential player's task is to announce a sequence of existential moves (i.e. choices about which child of an existential node in the computation tree should be pursued) without knowing in general what the universal player is doing for its part, and still end up in an accept state.

The tricky point in the PAFA definition is to overrule all computation subtrees which do not respect the privacy of the universal player, i.e. allowing the existential player to make different moves in response to two sequences of universal moves which should be identical from its point of view. For any existential node N in the computation tree, the *public moves seen by N* is defined to be the string in Γ^* corresponding to the list of public moves made in the path from the root down to N . Our additional criterion for an accepting computation subtree is that two existential nodes $N1$ and $N2$ with the same common state component can make different moves in this subtree only if the public moves seen by $N1$ and $N2$ are different. (Even if $N1$ and $N2$ are at different levels of this tree, they still have to satisfy this requirement to prevent existential choices being taken by “keeping time” in order to have knowledge about the position of the input head.)

A bonus of our practice of labeling the moves in the program becomes evident in the presentation of our PAFA algorithms. Recall that nondeterministic algorithms can be reformulated in a scenario where a deterministic verifier is presented a purported certificate of membership for the input string. Such a nondeterministic machine can be viewed as involving an existential player which is producing that certificate, i.e. the string of symbols corresponding to its choices of moves, as described above, trying to lead the machine to acceptance. In our algorithm descriptions, we will mean precisely that when we say that the machine *produces a string* by existential branchings. The difference with pure nondeterministic computation will be that universal branchings enable the certificate to be checked parallelly in several places in real time, and the privacy of the universal moves will make sure that the existential player has no opportunity to fool different universal branches by different responses, and instead has to provide a single certificate that has to pass all parallel tests to cause acceptance.

We will now present three theorems illustrating the power of private alternation in the case

of realtime finite automata.¹ The first theorem shows that, unlike realtime deterministic, nondeterministic, and alternating finite automata, PAFAs can recognize non-regular languages, even on unary alphabets.

Theorem 3.1 $\text{UPOWER} = \{1^m \mid m = 2^n \text{ for some } n \geq 0\}$ can be recognized by a PAFA.

Proof. We build a PAFA M . Given an input 1^m , M produces a string of length m over the alphabet $\{1, 1^+\}$ existentially as it goes over the input. M handles the empty string and the string with length one separately, by accepting only the latter. M expects the string produced to contain the symbol 1^+ as a marker instead of 1 at positions $\left\lfloor \frac{(2^i-1)m}{2^i} \right\rfloor$ for all $1 \leq i$. In other words, the existential player producing the string provides a certificate showing that the only divisor of the input length is 2. For this purpose, it divides the input into two by marking $\left\lfloor \frac{m}{2} \right\rfloor$ ($i = 1$). It also divides the second half into two by marking $\left\lfloor \frac{3m}{4} \right\rfloor$ ($i = 2$), and keeps going like this until it is not able to divide the remaining string into two equal parts. If the length of the remaining string is one, then $m = 2^n$ for some n . Otherwise, $1^m \notin \text{UPOWER}$.

M reads this certificate while reading the input. It branches universally on every marker 1^+ , to both check if the next marker is situated exactly at the midpoint of the remainder of the certificate, and to move on to the next marker, parallelly. It makes a universal branching in the beginning of the input to check the correctness of the position of the first marker, and move on to the second one, parallelly.

If it has not decided to check a marker for validity yet, M reads the k th character of the certificate while its head scans the k th symbol of the input. At the moment when it decides to check the validity of the i th marker, M 's input head should be scanning the $\left(\left\lfloor \frac{(2^i-1)m}{2^{i-1}} \right\rfloor\right)$ th symbol of the input, meaning $x = \left\lceil \frac{m}{2^{i-1}} \right\rceil$ symbols should remain to be scanned on the input. It will check whether the placement of the i th marker is correct. M is supposed to see the i th marker after scanning exactly $\left\lfloor \frac{x}{2} \right\rfloor$ more input symbols. To check this, M slows down the process of producing the certificate. It produces one symbol per every two input symbols scanned. If m is a perfect power, and the position of the i th marker is correct, $x = 2 \left\lfloor \frac{x}{2} \right\rfloor$, and M sees the i th marker when it is scanning the last input character. If the i th marker is not seen together with the last input symbol, either x is not divisible by 2, or the position of the marker is not correct. M rejects the input in this case. Hence, it is clear that M accepts an input if and only if the positions of all the markers are correct, and the string following the last marker is the string with length one, which implies m is a perfect power. \square

The next theorem demonstrates a technique which will turn out to be useful later.

Theorem 3.2 The language $\text{TWIN} = \{wcv \mid w, v \in \{0, 1\}^*\}$ can be recognized by a PAFA.

Proof. We construct a PAFA M . Given an input in the form of wcv , where $w, v \in \{0, 1\}^*$, M first makes a private universal branching on the left end-marker of the input. This produces the parallel threads that will compare the substring w or v , respectively, with the string that will be produced by existential branchings after this decision. The first branch starts making existential branchings into three different branches (corresponding to producing a string on

¹For a comprehensive study of the power of private alternation under various resource bounds, refer to [6].

$\{0, 1, c\}$) on each symbol while reading wc . At the same time, it checks if the string produced by the existential branchings is the same as wc . It accepts if they are the same, and rejects otherwise. The other branch first brings the input head on c , then it simply does similar three-way existential branchings on each character of v and the end-marker, comparing the 0's and 1's with the corresponding symbols of v , and the c with the end-marker. It accepts only if they all match.

If the input is of the form of wcw , there is a series of existential branches that produce the one string, wc , that works for both branches created by private universal branching in the beginning. Since the universal branching in the beginning is private, a set of existential branchings that makes the automaton accept in the end should be the same for both those branches created in private. In case of input wcv where $w \neq v$, there is no string z such that both $wc = z$ and $vc = z$. Therefore, the language of this machine is $\{wcw | w \in \{0, 1\}^*\}$. \square

More detailed proofs of Theorems 3.1 and 3.2 can be found in the Appendix.

Theorem 3.3 *The emptiness problem for PAFAs is undecidable.*

Proof. Assume that there exists a decision procedure for this problem. We can then decide whether a given TM M accepts an input w as follows: Construct a PAFA A , described below, which accepts a string only if it is an accepting computation history of M on w , (i.e. a string composed of segments representing the configurations that M would be in when started on w , all the way to an accepting configuration) properly encoded using 0's and 1's, use the emptiness test on A to see whether M accepts w , and announce the opposite of what the emptiness test says.

A branches existentially to produce a string which will be checked to see if it is the correct computation history of M on w starting from the second step (i.e. with the first segment supposed to represent the configuration resulting from the first move of M on w), and makes private-universal splits at the beginning of each segment of its input to 1) compare the i th segment it is reading with the $(i - 1)$ st one it is producing for equality (for $i > 1$), employing the method described in the previous theorem, and 2) process the i th segment it is reading to see if its successor according to the transition rules of M would equal the i th segment it is producing.

A also makes sure that the first segment of its input is indeed the start configuration of M on w , and that the final segment of its input contains the accept state in order to accept. \square

The realtime private alternating automaton with one counter (PA1CA) model is obtained by augmenting the PAFA with an integer counter that starts out at zero, and that can only be accessed by the universal states (only δ_U is modified accordingly). We now show that the language $\text{USQUARE} = \{1^m | m = n^2 \text{ for some } n \geq 0\}$ that is thought to be difficult for realtime counter automata (in fact, the only type of counter automaton known to recognize it until now is an exponential-time quantum version with two-way access to the input [23]) is indeed recognized when the counter automaton is augmented by private alternation.

Theorem 3.4 *USQUARE can be recognized by a PA1CA.*

Proof. We describe a PA1CA M for **USQUARE**. M will accept if it verifies that it can existentially produce a string whose length equals that of the input string, and which is the concatenation of x substrings (*segments*) of the form $1^{x-1}\#$ for some integer x . M starts by branching universally to perform two parallel controls on the produced string. For the first control, M uses private universal branchings to select any possible pair of segments, increments the counter for every x in the first segment of this pair, and decrements it for each x of the second segment, entering an accept state only if the counter returns to zero precisely when at the end of the second segment. The second control is on the number of segments provided in the certificate: The counter is incremented by one on every symbol but the final $\#$ of the first segment, and decremented by just one on each segment after that. This branch of M enters an accept state only if the counter returns to zero just at the end of this process. M makes sure the length of the produced string and the input are the same by producing one symbol for every input symbol that it scans.

M therefore accepts if and only if there exists an x such that the input length is x^2 . \square

It should be noted that the machine described in the proof above checks whether the counter is zero or not only at the very end of the computation, making it a *blind-counter* machine.

4. Quantum alternation

A quantum finite automaton (QFA) [9, 26] is a quintuple

$$M = (Q, \Sigma, \{\mathcal{E}_\sigma \mid \sigma \in \tilde{\Sigma}\}, |v_0\rangle, P),$$

where $Q = \{q_1, \dots, q_n\}$, $\mathcal{E}_\sigma = \{E_{\sigma,1}, \dots, E_{\sigma,l_\sigma}\}$ is the superoperator for the symbol $\sigma \in \tilde{\Sigma}$ composed by l_σ operation elements, $|v_0\rangle \in \{|q_1\rangle, \dots, |q_n\rangle\}$ is the initial state, and $P = \{P_a, P_r\}$ is the measurement operator applied after reading the right end-marker. An input is accepted if the outcome “a” of P is observed. For any given input $w \in \Sigma^*$, the computation of M can be traced by a $|Q| \times |Q|$ -dimensional density operator (mixed state)²:

$$\rho_j = \mathcal{E}_{w_i}(\rho_{j-1}) = \sum_{k=1}^l E_k \rho_{j-1} E_k^\dagger,$$

where $\rho_0 = |q_0\rangle\langle q_0|$ and $1 \leq j \leq |w|$, and the accepting probability of M on w is

$$f_M(w) = \text{Tr}(P_a \rho_{|w|}).$$

The nondeterministic QFA (NQFA) model can be defined using an acceptance mode known as *positive one-sided unbounded error* [1, 25]: A language L is said to be recognized by a NQFA N if each member of L is accepted by N with a positive probability, and each member of \bar{L} is accepted by N with zero probability.

²A pure state is a basis state or a linear combination of basis states with norm 1. A mixed state is a mixture of pure states such that the system is in each pure state with a nonzero probability and the summation of all probabilities is 1, i.e. $\{(|\psi_i, p_i\rangle \mid 1 \leq i \leq t \text{ and } \sum_{i=1}^t p_i = 1)\}$ for $t > 0$ different pure states.

Theorem 4.1 *The emptiness problem for NQFAs is decidable, where the automata are defined using algebraic numbers as transition amplitudes. Moreover, if the transitions amplitudes are restricted to rational numbers, we can also give a time bound.*

Proof. We can easily design a 1-state QFA with rational amplitudes, say Q , which accepts all input with probability 0. Let N be the NQFA we are supposed to test. Firstly, we assume that N is defined with rational numbers. The equivalence problem for QFAs, i.e. deciding whether any two given QFAs accept any given input with the same probability, is solvable in polynomial time if the automata are defined with rational amplitudes [22, 11]. Therefore, we can easily design a polynomial-time algorithm that takes Q and N as the input and then determines whether they are equivalent or not. If there are equivalent, then N defines an empty set, too.

Now, we assume that N is defined with algebraic numbers. The minimisation problem for QFAs, i.e. taking a QFA as the input and then outputting a minimal equivalent QFA, is also solvable if the amplitudes are restricted to algebraic numbers [12]. Therefore, we can design an algorithm that takes N as the input, constructs its equivalent minimal QFA, and accepts only if that minimal machine has a single state which is a reject state. \square


The universal QFA (UQFA) model, which can be considered as the “complement” of the NQFA, is defined using an acceptance mode known as *negative one-sided unbounded error* [25]. A language L is said to be recognized by a UQFA U if each member of L is accepted by U with probability 1, and each member of \bar{L} is accepted by N with probability less than 1. For a given QFA with rational amplitudes, say Q , the problem of whether there is a string accepted by Q with probability $p \in [0, 1]$ is undecidable.³ Therefore, the emptiness problem for UQFAs is undecidable.

Corollary 4.2 *The emptiness problem for UQFAs with rational amplitudes is undecidable.*

The class of languages recognized by NQFAs (resp., UQFAs) form a superset of the regular languages [25], called exclusive (resp., co-exclusive) stochastic languages (S^\neq (resp., $S^=$)) [14]. S^\neq and $S^=$ do not contain any nonregular *unary* languages [20]. Therefore, it is interesting to ask whether alternating between the capabilities of the two could lead us to recognizing some nonregular unary languages. It is known that such “two-alternation” is sufficient for recognizing the well-known NP-complete language SUBSETSUM (or KNAPSACK) [24]. So, it is expectable that two-alternation is sufficient for unary nonregular languages. In the following, we will first describe a model of quantum alternation that embodies these ideas, and then present a result for USQUARE, using the methods introduced in [24].

An alternating quantum finite automaton (AQFA) has both classical and quantum states [24]. The classical states are once again either existential or universal, with some of them designated as accepting states. The computation is controlled by the classical states. Each step has two phases. In the first phase (quantum phase), depending on the current symbol and classical state, a superoperator is applied to the quantum state, and an outcome having non-zero probability is observed. Figure 1 represents the outcomes of a superoperator applied to a pure quantum

³Note that this result can be obtained even for the known simplest QFA model [13], known as measure-once or Moore-Crutchfield QFA [7, 3].


 Figure 1: The pure states after applying a superoperator with l operation elements

state, where indices of the operators are the outcomes. Note that the outcomes are discarded in the above definition of QFA, but they are still important for AQFAs, since the branches of the computation are defined based on them. In the second phase (classical phase), the new classical state is determined by the current symbol, the classical state, and the outcome obtained in the first phase. Similar to the previous models, we assume that the automaton spends two steps on each symbol. The computation of an AQFA can be shown as a tree such that each node represents a pair consisting of a classical state and a pure quantum state, and the edges represent the transitions with non-zero probabilities. Note that the exact magnitudes of these probabilities are not significant. We associate each inner node with either a “ \vee ” or a “ \wedge ”, depending on whether its corresponding classical state is existential or universal, respectively. When evaluating the tree, each leaf in which the input is accepted (rejected) takes the value of “true” (“false”). These halting decisions are given by checking whether the computation ends in a classical accepting state. The value of root can be evaluated from bottom to top: Any existential inner node (assigned a “ \vee ”) takes the value of “true” if at least one of its children has taken the value of “true” and any universal inner node (assigned a “ \wedge ”) takes the value of “true” if all of its children have taken the value of “true”. The input is accepted if and only if the root takes the value of “true”. Note that classical alternation is a special case of quantum alternation, and NQFAs and UQFAs are just AQFAs with only existential and universal states, respectively.

Theorem 4.3 *The language USQUARE can be recognized by a AQFA with two-alternation.*

Proof. We use the encoding techniques given in [24]. Let $w = a^m$ be the given input, where $m > 0$. The idea is as follows: The automaton nondeterministically picks a position on w , say i , and encodes the values of i^2 and m into the amplitudes of two quantum states. Technically, these two values are kept in a pure state together with some auxiliary ones, and each value is divided by a coefficient so that the pure state has norm 1. At the end of the input, the only universal choice is made: The input is accepted in all the branches except the one in which the amplitudes keeping i^2 and m are subtracted from each other and the input is rejected based on the result amplitude. It is clear that if w is a member of USQUARE, then this branch does not appear on the computation tree for $i = \sqrt{m}$, and so the corresponding universal node takes the value of “true” and the value of the root becomes “true”, too. On the other hand, if w is not

a member, then, for any choice of i , this rejecting branch always appears on the computation tree, and so the corresponding universal nodes take the value of “false” and the root cannot take the value of “true”. \square

We now show that the emptiness problem for AQFAs over unary alphabets is undecidable.

Theorem 4.4 *The emptiness problem for AQFAs over unary alphabets is undecidable.*

Proof. It is known that one-way AQFAs can simulate the computation of a given TM on any input [24]. In that simulation, the automaton first reads the whole input, and the head arrives at the right end-marker. From that point on, the head never moves, and it nondeterministically selects the next configurations, with each wrong guess is eliminated by universal states. When a halting configuration appears, the automaton gives a parallel decision. The automaton runs forever but if the TM halts on the input, then the decision can be derived from a finite level of the computation tree.

Let M be any TM, w be an input, and O be a one-way AQFA simulating the computation of M on any given input. For M and input w , we can design a (realtime) AQFA, say R , that execute some finite steps (depending on the length of the input) of the simulation implemented by O for (M, w) . Here, R can execute more steps of the simulation on the longer input string and there is no difference between binary and unary input strings, i.e. only the input length is important. If M halts on w , then there are some (actually infinitely many) sufficiently long unary strings such that the realtime AQFA can give the decision of M on w on these unary inputs.

Therefore, if the emptiness problem for realtime AQFAs with unary alphabet is decidable, then the halting problem for TMs is decidable, too. But, this is a contradiction and so the emptiness problem is undecidable. \square

Acknowledgements. We thank the anonymous reviewers for their helpful comments.

References

- [1] L. M. ADLEMAN, J. DEMARRAIS, M.-D. A. HUANG, Quantum Computability. *SIAM Journal on Computing* 26 (1997) 5, 1524–1540.
- [2] M. D. BIASI, A. YAKARYILMAZ, Non-trivial unary languages recognized by two-way one-counter machines. In: *CIAA. LNCS 8587*, Springer, 2014, 148–161. (arXiv:1404.2761).
- [3] V. D. BLONDEL, E. JEANDEL, P. KOIRAN, N. PORTIER, Decidable and Undecidable Problems about Quantum Automata. *SIAM Journal on Computing* 34 (2005) 6, 1464–1473.
- [4] A. K. CHANDRA, D. C. KOZEN, L. J. STOCKMEYER, Alternation. *Journal of the ACM* 28 (1981) 1, 114–133.
- [5] A. CONDON, *Computational Models of Games*. MIT Press, 1989.
- [6] H. G. DEMIRCI, A. C. C. SAY, A. YAKARYILMAZ, The complexity of debate checking. *Theory of Computing Systems* (2014). DOI:10.1007/s00224-014-9547-7.

- [7] H. DERKSEN, E. JEANDEL, P. KOIRAN, Quantum automata and algebraic groups. *Journal of Symbolic Computation* 39 (2005) 3–4, 357–371.
- [8] S. GINSBURG, H. G. RICE, Two Families of Languages Related to ALGOL. *Journal of the ACM* 9 (1962) 3, 350–371.
- [9] M. HIRVENSAALO, Quantum Automata with Open Time Evolution. *International Journal of Natural Computing* 1 (2010) 1, 70–85.
- [10] O. H. IBARRA, T. JIANG, N. Q. TRÂN, H. WANG, On the Equivalence of Two-Way Pushdown Automata and Counter Machines Over Bounded Languages. *International Journal of Foundations of Computer Science* 4 (1993) 2, 135–146.
- [11] L. LI, D. QIU, X. ZOU, L. LI, L. WU, P. MATEUS, Characterizations of one-way general quantum finite automata. *Theoretical Computer Science* 419 (2012), 73–91.
- [12] P. MATEUS, D. QIU, L. LI, On the complexity of minimizing probabilistic and quantum automata. *Information and Computation* 218 (2012), 36–53.
- [13] C. MOORE, J. P. CRUTCHFIELD, Quantum automata and quantum grammars. *Theoretical Computer Science* 237 (2000) 1-2, 275–306.
- [14] A. PAZ, *Introduction to Probabilistic Automata*. Academic Press, New York, 1971.
- [15] H. PETERSEN, Two-way one-counter automata accepting bounded languages. *SIGACT News* 25 (1994) 3, 102–105.
- [16] G. L. PETERSON, J. H. REIF, Multiple-person alternation. In: *SFCS’79: Proceedings of the 20th Annual Symposium on Foundations of Computer Science*. 1979, 348–363.
- [17] J. H. REIF, Universal Games of Incomplete Information. In: *Proceedings of the 11th Annual ACM Symposium on Theory of Computing*. STOC’79, ACM, 1979, 288–308.
- [18] J. H. REIF, The Complexity of Two-Player Games of Incomplete Information. *Journal of Computer and System Sciences* 29 (1984) 2, 274–301.
- [19] K. REINHARDT, A. YAKARYILMAZ, The minimum amount of useful space: New results and new directions. In: *Developments in Language Theory*. LNCS 8633, Springer, 2014, 315–326. (arXiv:1405.2892).
- [20] A. SALOMAAA, M. SOITTOLA, *Automata-Theoretic Aspects of Formal Power Series*. Texts and monographs in computer science, Springer-Verlag (New York), 1978.
- [21] M. SIPSER, *Introduction to the Theory of Computation, 3rd edition*. Cengage Learning, United States of America, 2013.
- [22] W.-G. TZENG, A Polynomial-Time Algorithm for the Equivalence of Probabilistic Automata. *SIAM Journal on Computing* 21 (1992) 2, 216–227.
- [23] A. YAKARYILMAZ, One-Counter Verifiers for Decidable Languages. In: *CSR*. LNCS 7913, Springer, 2013, 366–377.
- [24] A. YAKARYILMAZ, Quantum Alternation. In: *CSR*. Lecture Notes in Computer Science 7913, Springer, 2013, 334–346.
- [25] A. YAKARYILMAZ, A. C. C. SAY, Languages recognized by nondeterministic quantum finite automata. *Quantum Information and Computation* 10 (2010) 9&10, 747–770.
- [26] A. YAKARYILMAZ, A. C. C. SAY, Unbounded-error Quantum Computation with Small Space Bounds. *Information and Computation* 279 (2011) 6, 873–892.

1. Detailed proofs of Theorems 3.1 and 3.2

Theorem 3.1 $\text{UPOWER} = \{1^m | m = 2^n \text{ for some } n \geq 0\}$ can be recognized by a PAFA.

Proof. We build a PAFA M . M has a specific existential state (e, α) where e and α are the common and private state components, respectively. (e, α) has only two outgoing transitions to two different universal states: (u_0, α) and (u_1, α) (remember that only the common state component can be changed via a transition from an existential state). (u_0, α) has only one outgoing transition: back to (e, α) . (u_1, α) has one transition back to (e, α) and another transition to (e, β) when reading an a on the input. Similarly, (e, β) has transitions to (u_0, β) and (u_1, β) . Note that (e, β) has the same common component with (e, α) , so they have to have the same exact sequence of choices throughout the computation between changing the common state component to either u_0 or u_1 . This remark is key to understanding the power of private alternation. (u_0, β) has only one transition when reading an a on the input, that is, to the existential state (e', β) which goes to the universal state (u', β) which goes back to (e, β) on input symbol a . (u_1, β) goes to an accepting state if it reads the end marker $\$$, and goes to a rejecting state if it reads an a . All the other universal states described above also make a transition to an accepting state upon reading the end marker $\$$.

Although the operation of the main part of M described above may not be obvious at first, it is easier to understand when the following fact is realized. The existential player has a choice to make different transitions (to change the common state component to either u_0 or u_1) only on states (e, α) and (e, β) which have the same common state component e . Since the choice existential player makes cannot depend on anything but the common state component which is identical for both (e, α) and (e, β) , this means that the existence of a predetermined sequence of choices for existential branchings which would lead to acceptance no matter what universal choices are realized is the only condition for M to accept an input.

Now we explain the part carried out by universal states. In the part where the private state component is α , the only task (u_0, α) performs is to make a transition back to (e, α) and let existential player make another choice between u_0 and u_1 . Note that the input head moves forward for each such two transitions (i.e. (e, α) to (u_0, α) , and (u_0, α) to (e, α) back again). However, every time existential player makes the choice to make a transition to (u_1, α) , (u_1, α) creates two different branches one leading to (e, α) and letting it continue as it was, and the other transferring the control to (e, β) . The part with the private state component β carries out a control over the choices existential player makes. To simply put it, it checks if the number of u_0 choices made by the existential player until the very next u_1 choice is exactly half the number of remaining input symbols on the input tape (i.e. number of symbols until the right end marker). To implement this, M has two dummy states (e', β) and (u', β) which makes M read two input symbols per each u_0 choice made by the existential player.

On input a^m , consider the following strategy for the existential player. He will make the choice to replace the common state component with u_0 on every choice except his $\left\lfloor \frac{(2^i - 1)m}{2^i} \right\rfloor$ th choice for all $1 \leq i$ where he replaces it with u_1 . In other words, the first u_1 choice corresponds to $\left\lfloor \frac{m}{2} \right\rfloor$ th choice ($i = 1$) dividing m by 2. Second such choice corresponds to $\left\lfloor \frac{3m}{4} \right\rfloor$ ($i = 2$) dividing

the remaining $\frac{m}{2}$ input symbols after the first choice by 2, and keeps going like this until it is not able to divide the remaining string into two equal parts. If the length of the remaining string is one, then $m = 2^n$ for some n . Otherwise, $a^m \notin \text{UPOWER}$. Since the only control carried out by the universal states (in part β) is to, for every u_1 choice, check if the number of u_0 choices before the very next u_1 choice is exactly half the number of remaining input symbols, this strategy will not lead to any rejecting leaf. Therefore, M accepts if $m = 2^n$ for some $n \geq 0$. Furthermore, due to the nature of control carried out by the universal states in part β , this is the only strategy (to keep dividing the remaining part of the input into two in this manner) that would lead to acceptance.

The part of M , described above does not work for empty string and the string with length 1, but a couple of additional states can be integrated to handle these two cases separately. \square

Theorem 3.2 *The language $\text{TWIN} = \{wcw | w \in \{0, 1\}^*\}$ can be recognized by a PAFA.*

Proof. We will first construct a PAFA M by giving the states and the transitions of the main part of M and, then, give the intuition behind it. M has the initial state (e_i, p) with the only outgoing transition to (u_i, p) where two different universal branches are created upon reading the left end marker: one to (e, α) , and the other to (u_i, β_1) . (e, α) has an existential choice between making a transition to (u_0, α) or to (u_1, α) . (u_0, α) goes to (e, α) upon reading a 0 on input, and it goes to a rejecting state upon reading a 1. Conversely, (u_1, α) makes a transition to (e, α) upon reading a 1 on input, and it goes to a rejecting state if it reads a 0. Both (u_0, α) and (u_1, α) ends up in an accepting state if they read a c or the end marker $\$$ on the input.

The task carried out by (u_i, β_1) , on the other hand, is fairly simple. Every time, it makes a transition to (u_i, β_2) which goes back to (u_i, β_1) upon reading a 0 or a 1, and makes a transition to (e, α) upon reading c . Basically, what these two states with private components β_1 and β_2 do is to bring the input head on the character right after c and transfer the control to the part with private state component α described above.

In the part with the private state component α , M makes a new existential choice on each input symbol using state (e, α) . This choice is between making a transition to (u_0, α) and making a transition to (u_1, α) . It keeps making such choices as long as its choices on 0's are to (u_0, α) , and on 1's are to (u_1, α) . Otherwise, (u_0, α) and (u_1, α) will end the computation by making a transition to a rejecting state. A more plausible way of interpreting these existential choices is to think of a transition to (u_0, α) as the symbol 0 and a transition to (u_1, α) as the symbol 1, and treat the sequence of choices as a certificate string from $\{0, 1\}^*$. In this way, as long as this *string produced by* the existential player is the same with the string read on the input, (u_0, α) and (u_1, α) never cause M to reject.

On input wcv , the universal branch from (u_i, p) to (e, α) makes M compare the string produced by the existential choices with w (the part before c). Since the universal branch from (u_i, p) to (u_i, β_1) first brings the input head on the first character of v and, then, transfers control to the part with private component α , this branch makes M also compare the the string produced by the existential choices with v . M rejects if and only if one of these comparisons fails. The strings produced by the existential choices for these two branches are bound to be the same due to the definition of PAFA. In other words, existential player cannot differentiate between transitions

from (u_i, p) and from (u_i, β_2) because the common state components seen by the existential player up to that point is identical (e_i and u_i) for both transitions. Hence, M compares an existentially produced string with both w and v and accepts if and only if it is identical to both of them. This implies that M recognizes TWIN. \square